LLM + Neutrosophic fsQCA: Inter-Narrative Causal Consistency and Paraconsistent Detection in Media Accounts of Urban Violence in Guayaquil

Maikel Yelandi Leyva-Vazquez[1,2,3] | **Noel Batista Hernandez**[1] | **Florentin Smarandache**[4]

1. Universidad Bolivariana del Ecuador, Guayaquil, Ecuador

2. Universidad de Guayaquil, Guayaquil, Ecuador

3. Universidad Bernardo O'Higgins, Santiago, Chile

4. University of New Mexico, Gallup, NM, USA

Corresponding author: mleyvaz@gmail.com | ORCID: 0000-0001-7911-5879

PREPRINT — submitted to OSF Preprints — May 2026 — Not peer reviewed

Abstract

Urban violence in Guayaquil, Ecuador has reached crisis levels, yet causal explanations remain fragmented across local and international media. This paper introduces N-fsQCA, a Neutrosophic extension of fuzzy-set Qualitative Comparative Analysis (fsQCA), and combines it with Large Language Models (LLMs) to extract and compare causal narratives from a bilingual corpus of 31 sources (16 local Spanish-language, 15 international English-language). Four LLMs (Google Gemini-2.0-Flash-Lite, Meta LLaMA-3.1-8B, Microsoft Phi-4, Qwen-3-8B) assign fuzzy scores [0,1] to eight causal conditions; indeterminacy (I) is operationalized as inter-LLM variance (Var/0.25), capturing epistemic disagreement as structural information. Results show: (i) the configuration territorial_war * prison_linkages achieves the highest consistency (T=0.908, I=0.092); (ii) international media emphasizes drug routes (gap=+0.330) and prison linkages (gap=+0.309) at roughly double the rate of local media; (iii) five structural drivers are systematically silenced in the press, with weapons trafficking showing the largest gap (-0.444) against a structural baseline derived from a validated perception survey (n=179) and grey literature. The pipeline is open-source (MIT License).

Keywords: Neutrosophic fsQCA · LLM · urban violence · media narratives · causal consistency · Guayaquil · paraconsistency · inter-narrative analysis

1. Introduction

Urban violence in Guayaquil, Ecuador represents one of the most dramatic security deteriorations in Latin America. After a 14% fall in 2023-2024, homicide projections for 2025 estimate around 9,100 deaths — a 40% increase — driven by criminal network fragmentation following the arrest and extradition of major gang leaders (Global Initiative, 2025; Crisis Group, 2025; ACLED, 2025). Despite extensive coverage, explanations diverge sharply across media contexts: local Spanish-language media foregrounds gang fragmentation and extortion; international English-language media emphasizes drug-trafficking routes and prison coordination.

Ecuador's security crisis reflects structural conditions that intersect international drug-trafficking networks, collapsing prison governance, and decades of socioeconomic exclusion in coastal urban areas. The port city of Guayaquil, which handles nearly 70% of the country's external trade, has become a critical node in transnational cocaine flows from Colombia and Peru to European markets, attracting violent competition among criminal organizations that previously operated under relative cartel discipline (InSight Crime, 2025a; Crisis Group, 2025). The disarticulation of gang hierarchies following the arrest of key leaders unleashed fragmented inter-band warfare that proved harder to contain than the preceding era of centralized control (ACLED, 2025). Against this backdrop, media coverage has multiplied, yet its capacity to accurately represent causal structure remains largely unexamined. Homicides are reported; their determinants are narrated selectively, with local and international outlets encoding structurally different explanatory theories of the same events.

This divergence raises an empirically tractable question: do different media frames constitute genuinely different causal theories about the same phenomenon, or do they reflect structural asymmetries in journalistic attention? Answering it requires a method that can (a) extract causal conditions from unstructured text at scale, (b) evaluate configurational consistency across sources, and (c) quantify epistemic disagreement without collapsing it into a single binary judgment.

The challenge is methodological as much as empirical. Existing approaches to media framing analysis rely on manual coding, which does not scale, or topic models, which recover themes but not causal structure. Simple LLM classification produces a single point estimate without quantifying disagreement across models or sources. None of these approaches can simultaneously extract causal configurations from unstructured bilingual text, assess whether independently extracted explanations converge on the same configurational logic, and quantify the epistemic uncertainty inherent in that extraction process. Addressing all three requirements is the technical motivation for the N-fsQCA pipeline introduced in this paper.

This paper provides that method. We combine LLMs for fuzzy causal scoring with N-fsQCA, a Neutrosophic extension of fsQCA (Ragin, 2008; Smarandache, 1998) that represents indeterminacy (I) as a first-class component alongside truth (T) and falsity (F).

Indeterminacy is operationalized as inter-LLM variance — the disagreement among four language models becomes a measure of epistemic uncertainty, not noise.

The remainder of the paper is organized as follows. Section 2 reviews related work on LLM causal extraction, configurational analysis, neutrosophic logic, and Latin American media framing. Section 3 develops the theoretical framework, formalizing inter-narrative causal consistency and the N-fsQCA triplet. Section 4 describes the methodology: corpus construction, causal condition operationalization, LLM scoring protocol, structural baseline derivation, and N-fsQCA execution. Section 5 presents results across five analyses: corpus characterization, dominant causal configurations, LLM model ablation, local-versus-international framing comparison, and silencing analysis. Section 6 discusses implications, Section 7 acknowledges limitations, and Section 8 concludes.

1.1 Contributions

This paper makes five contributions. First, it introduces N-fsQCA, the first neutrosophic extension of fsQCA, with indeterminacy operationalized as I = Var(LLM_scores) / 0.25. Second, it proposes inter-narrative causal consistency as an evaluation criterion for LLM causal explanations in domains without ground truth. Third, it provides an empirical analysis of 31 bilingual sources (16 ES / 15 EN) on Guayaquil's violence crisis, revealing systematic local/international framing asymmetries. Fourth, it performs a quantified silencing analysis using a validated perception survey (n = 179) as a structural baseline. Fifth, the pipeline is fully open-source and replicable at under USD 5 (github.com/mleyvaz/llm-nfsqca-narrative-violence).

2. Related Work

2.1 LLM causal extraction

LLMs have demonstrated strong performance in causal extraction tasks (Tonmoy et al., 2024). Evaluation methods include semantic entropy (Farquhar et al., 2024) and consistency-based hallucination detection (Manakul et al., 2023). No published work applies these to configurational narrative analysis.

A persistent limitation of LLM-based causal extraction is the absence of a domain-agnostic evaluation criterion for settings where empirical ground truth does not exist. In tasks with labeled corpora such as clinical records or engineering incident reports, precision and recall against gold annotations provide a clear target. In open domains such as conflict journalism, no authoritative causal annotation exists, and the correct causal structure is itself contested among domain experts. Existing alternatives include semantic entropy (Farquhar et al., 2024), which quantifies distributional uncertainty over generated explanations, and self-consistency checking via paraphrase (Manakul et al., 2023), which detects contradictions within a single model's outputs. Neither approach allows comparison of causal

configurations extracted independently by multiple models from multiple sources. Inter-narrative consistency evaluated across a corpus fills that gap.

2.2 fsQCA and configurational causality

fsQCA (Ragin, 1987, 2000, 2008) operationalizes set-theoretic causality through Boolean minimization of truth tables. Recent work combines QCA with automated text processing but without neutrosophic extension.

The key insight of set-theoretic causality, as opposed to correlational approaches, is that social outcomes are typically produced by conjunctural causation: it is the combination of conditions, not their independent additive contributions, that matters (Misangyi et al., 2017). Territorial war alone may not produce a violence spike; territorial war combined with weakened prison governance and disrupted drug routes may be sufficient. fsQCA formalizes this through Boolean minimization of a truth table, yielding configurations that are sufficient for the outcome. When applied to media narratives, configurations represent journalistic causal theories: the implicit explanatory logic embedded in how sources narrate an event. Comparing configurations across source types thus operationalizes framing divergence at the level of causal structure rather than mere topic co-occurrence. No published work, to our knowledge, has combined LLMs with fsQCA using a neutrosophic extension that allows indeterminacy to be represented as a first-class component.

2.3 Neutrosophic logic

Neutrosophic logic (Smarandache, 1998, 2003) generalizes fuzzy sets by adding I (indeterminacy), enabling paraconsistent representation where $T + F > 1$. Applications include multi-criteria decision making, machine learning, and educational assessment (Batista Hernandez et al., 2023). N-fsQCA is the first application to configurational causal analysis of media discourse.

The paraconsistent regime, where $T + F > 1$, is particularly important for narrative analysis. In classical and fuzzy logic, a proposition cannot be simultaneously affirmed and denied above a half-truth threshold. In neutrosophic logic, following the tradition of paraconsistent logics (Da Costa, 1974), $T + F > 1$ is a legitimate epistemic state indicating that genuinely contradictory evidence has been documented. Applied to media corpora, this regime captures the situation where a causal condition is strongly asserted by some sources and strongly absent or denied in others, not because one side is wrong, but because sources encode structurally different theories. The ability to detect this paraconsistent state without forcing a binary resolution is a core advantage of N-fsQCA over classical fsQCA. Single-valued neutrosophic sets (SVNS; Wang et al., 2010) provide the algebraic foundation, with the constraint $T + I + F <= 1$ enforced in the standard regime and $T + I + F > 1$ flagging paraconsistency.

2.4 Media framing of Latin American violence

Critical media studies document systematic divergences in violence framing across media contexts (Mhlambi, 2020; Birhane, 2021; Van Dijk, 2006). This paper provides the first quantitative configurational operationalization of such divergences.

The decolonial media studies literature argues that international coverage of Latin American violence systematically privileges geopolitical and supply-chain framings over the structural economic conditions that community actors identify as primary drivers (Mhlambi, 2020; Birhane, 2021). Van Dijk (2006) and Fairclough (1995) demonstrate how securitization frames exported via wire services systematically displace local narratives that foreground structural inequality and institutional abandonment. The present paper operationalizes this theoretical claim through a quantitative configurational analysis that makes the framing asymmetry measurable and falsifiable. By grounding the comparison in a structural baseline derived from a community perception survey, the analysis connects media representations directly to the causal understanding of affected populations, a methodological step that previous quantitative framing studies have not taken.

2.5 Discourse pragmatics and causal silence

The analysis of what press accounts omit, as much as what they assert, requires grounding in the pragmatics of language. Grice's (1975) cooperative principle holds that communicators implicitly commit to making their contributions as informative as required for the current purposes of the exchange (maxim of quantity) and to saying only what they have adequate evidence for (maxim of quality). Applied to causal journalism, systematic omission of a driver that community actors recognize as salient constitutes a violation of the quantity maxim: the text is less informative than the situation requires. This violation is not incidental but structural, reproduced across dozens of independent sources, which is what the silencing analysis in Section 5.5 measures. Relevance theory (Sperber and Wilson, 1995) deepens the account: what gets communicated as causally relevant in a news article reflects the producer's model of the assumed cognitive environment of their intended audience. A drug-trafficking frame is relevant to international audiences with a supply-chain mental model; an extortion-and-territorial-control frame is relevant to local audiences whose daily life intersects the extortion economy directly. Neither frame is informationally deficient in isolation; the divergence becomes visible only through cross-context comparison of the kind that N-fsQCA operationalizes.

Critical Discourse Analysis (CDA) operationalizes these insights for news texts specifically. Van Dijk (2006) demonstrates that the selection of causal agents and processes in crime journalism is ideologically structured: elite sources are more likely to be cited as causal authorities, structural conditions are more likely to be backgrounded as presupposed context rather than foregrounded as asserted causes, and oppositional or community-based causal theories are systematically marginalised. Fairclough (1995) extends this to the inter-textual dimension: news discourse is produced within genre conventions that constrain which causal frames are recognisable as journalistically legitimate. These conventions differ across media systems, which is precisely why local Ecuadorian and international English-

language outlets encode structurally different causal theories of the same events. Connecting CDA to N-fsQCA yields a concrete operational hypothesis: conditions that are structurally backgrounded in a media system will produce low T values in LLM extraction, because the causal attribution is expressed through presupposition and implicature rather than direct assertion, and will therefore be harder for extraction systems to score as salient. The silencing gaps in Section 5.5 test this hypothesis quantitatively.

3. Theoretical Framework

3.1 Inter-narrative causal consistency

In domains without empirical ground truth, we define inter-narrative causal consistency as the degree to which independently extracted LLM explanations converge on the same configurational causal structure. This criterion shifts evaluation from 'is the claim factually correct?' to 'do independent extraction processes agree on the causal structure?'

This criterion has several desirable properties. First, it is scalable: because it relies on automated LLM extraction rather than manual annotation, it can be applied to large multilingual corpora without incurring prohibitive annotation costs. Second, it is model-agnostic: the consistency is computed across an ensemble of heterogeneous LLMs, meaning that the criterion does not privilege any single model's idiosyncratic tendencies. Third, it is interpretable in set-theoretic terms: consistency values map directly onto fsQCA sufficiency scores, providing a formal connection between narrative agreement and configurational causality. Fourth, and most importantly, it is falsifiable: a low inter-narrative consistency score on a given configuration constitutes evidence that the journalistic causal theory encoded in different source types is genuinely divergent, not merely noisy. This falsifiability distinguishes N-fsQCA from purely qualitative accounts of framing divergence.

Grounding inter-narrative causal consistency in pragmatic theory adds a second, complementary interpretation. In discourse coherence terms (Hobbs, 1990; Mann and Thompson, 1988), a narrative is internally coherent when its propositions stand in recognised rhetorical relations: cause, result, elaboration, contrast. Inter-narrative causal consistency, as defined here, measures whether independently produced narratives from different sources assign the same cause-result structure to the same outcome. High consistency (high T) means that the causal rhetoric is shared across the corpus: different journalists, writing for different audiences, arrive at the same configurational logic. Low consistency combined with high indeterminacy (high I) signals a contested causal rhetoric: the rhetorical relation between a condition and the violence outcome is unstable across sources, reflecting genuine disagreement about the causal architecture of the crisis. This distinction between shared causal rhetoric and contested causal rhetoric is not available in approaches that collapse all sources into a single aggregate score; it requires the configurational apparatus of fsQCA extended with the indeterminacy component of neutrosophic logic.

The silencing phenomenon measured in Section 5.5 can be formalised as a corpus-level violation of the Gricean quantity maxim. Let B_k denote the structural truth value of condition k (the baseline) and P_k denote the press truth value (mean LLM score across all sources). A silencing event is defined as $P_k \ll B_k$, meaning that the press, taken as a collective communicator, systematically asserts condition k at a lower degree of salience than the structural record warrants. Under the quantity maxim, a cooperative communicator would make the causal attribution as salient as the evidence justifies; the gap $B_k - P_k$ therefore measures the degree of quantity-maxim violation at the condition level. This formalisation connects the quantitative result in Table 4 to a theoretical claim in pragmatic linguistics: press coverage of urban violence in Guayaquil is systematically less informative than community knowledge and expert evidence would require, and the silencing is differentially distributed across conditions in ways that reflect both practical constraints on journalism (weapons trafficking is politically sensitive) and structural features of global media systems (international outlets have different relevance models than local ones).

3.2 N-fsQCA components

For causal configuration C and outcome Y, the neutrosophic triplet (T, I, F) is defined as follows.

The three components jointly characterize the epistemic status of each configuration. High T and low I indicates a robust causal narrative: multiple LLMs, processing multiple sources, converge on the configuration as sufficient for the outcome. High T combined with high I indicates that the configuration achieves consistency but is disputed across models, suggesting that the causal claim is present but contested. High F alone would indicate a configuration that multiple sources explicitly contradict as a cause of the outcome. The paraconsistent regime, T + F > 1, arises when some sources strongly endorse and others strongly deny a configuration, signaling a genuine narrative fracture in the corpus. In practice, the Guayaquil corpus does not produce T + F > 1 because all sources agree on the outcome (violence spike), and F is therefore universally zero; this structural property of the corpus is acknowledged as a limitation in Section 7 and motivates the design of future studies that include low-violence comparison periods.

T (Truth / sufficiency): fuzzy consistency of C => Y, computed as $sum\big(min(X_T, Y_T)\big)/sum(X_T)$ following Ragin (2008), where X_T = mean LLM score across models.

I (Indeterminacy): Var(LLM_scores) / 0.25, operationalizing inter-LLM disagreement as epistemic uncertainty. Maximum variance for [0,1] values is 0.25.

F (Falsity): proportion of cases where $X_T > 0.5$ and $Y_T < 0.5$ (contradictory cases).

Normalization enforces T + I + F <= 1 (SVN constraint, Wang et al., 2010). The archetype STRONG_SUFFICIENT requires T >= 0.75 and I <= 0.25.

The STRONG_SUFFICIENT archetype serves as the primary classification threshold throughout the analysis. A configuration classified as STRONG_SUFFICIENT carries both empirical and communicative significance: empirically, it means that sources presenting the configuration as causal overwhelmingly also present the outcome; communicatively, it means that the causal claim is stable enough that readers of those sources would reliably form the same configurational expectation. Configurations that achieve T >= 0.75 but with I > 0.25 are classified as INCONSISTENT_UNCERTAIN, indicating that the causal narrative is present but that model disagreement is high enough to counsel caution in interpretation. Configurations with T < 0.75 and F = 0 are classified as WEAK, reflecting conditions mentioned inconsistently across the corpus without achieving systematic causal attribution.

4. Methodology

4.1 Corpus construction

We constructed a bilingual corpus of 31 sources with extractable text: 16 local Spanish-language sources (El Universo, Primicias, Vistazo, Extra, El Comercio, Infobae, France24-ES, SwissInfo-ES, OECO, EcuadorChequea, Diario San Rafael, El Comun, SwissInfo-2) and 15 international English-language sources (Global Initiative, ACLED, HRW, Al Jazeera, InSight Crime x3, CSIS, Atlantic Council, CBS News, FIU Gordon, Brookings, USCRI, AP News, Wikipedia-EN). Text was extracted via Trafilatura with browser-agent fallback. Sources span 2024-2026.

A clarification on the local/international classification is warranted. The distinction does not follow institutional origin but rather linguistic-audiential frame: sources are classified as Spanish-language (local) if they produce content in Spanish addressed to Latin American audiences using the journalistic conventions of the regional Spanish-language press, regardless of where the producing organization is headquartered. Under this criterion, Infobae (an Argentine digital outlet widely consumed in Ecuador), France 24 ES and SwissInfo ES (international broadcasters whose Spanish services target the Latin American Spanish-speaking public) are grouped with Ecuadorian sources because they share the same linguistic community of reception and draw on the same repertoire of regional journalistic frames. What distinguishes the Spanish-language group from the English-language group is not geography but the communicative contract implied by the language choice: Spanish-language coverage addresses readers whose proximity to the events is lived and embodied, while English-language coverage addresses readers whose engagement is analytical and geopolitically distanced. This distinction is theoretically motivated by relevance theory (Sperber and Wilson, 1995): the assumed cognitive environment of a reader consuming El Universo or Infobae differs systematically from that of a reader consuming InSight Crime or the Atlantic Council, and that difference in assumed cognitive environment constrains which causal frames are made maximally relevant in each outlet. The classification thus operationalizes a sociolinguistically grounded hypothesis rather than a geographic one.

Source selection followed three criteria. First, direct relevance: each source had to address the post-2022 violence escalation in Guayaquil specifically, not Ecuador or Latin American violence in general. Second, extractability: full-text extraction via Trafilatura was required, with a browser-agent fallback for paywalled or JavaScript-rendered pages; sources that yielded fewer than 300 words of clean text were excluded. Third, linguistic balance: the corpus was designed to achieve approximate parity between Spanish-language (local) and English-language (international) sources to permit cross-context comparison. The final corpus of 31 sources spans the period 2024-2026, covering both the peak violence phase and the initial response period following President Noboa's declaration of internal armed conflict in January 2024. Source diversity was prioritized across institutional type: legacy news organizations, investigative journalism outlets, think tanks, UN system reports, human rights organizations, and academic grey literature are all represented. No social media sources were included, as the extraction and attribution chain for informal content introduces uncontrollable noise into a corpus designed for systematic causal analysis.

4.2 Causal conditions

Code	Condition	Description
A	Gang fragmentation	Fragmentation/alliances among criminal bands
B	Drug routes/ports	Control of trafficking routes and port access
C	Extortion/vacunas	Protection payments and extortion economy
D	Territorial war	Disputes over zone control
E	Weapons trafficking	Firearms availability
F	Prison linkages	Coordination from penitentiaries
G	Institutional weakness	State corruption and governance failure
H	Economic pressure	Poverty, unemployment, income inequality

4.3 LLM scoring

Each of 4 LLMs (Google Gemini-2.0-Flash-Lite, Meta LLaMA-3.1-8B, Microsoft Phi-4, Qwen-3-8B) scored each condition [0,1] per article via structured prompt (JSON output, five-point anchored scale). T = mean across LLMs; I = Var / 0.25. One model (Mistral-7B v0.1) was excluded after failing JSON output in 7/31 cases, constituting an empirical finding on model reliability.

The scoring prompt was structured as follows. Each model was provided with the full extracted text of a source, a description of each causal condition (matching Table in Section 4.2), and a five-point anchored scale: 0.0 (condition completely absent or explicitly denied), 0.25 (condition briefly mentioned without causal attribution), 0.50 (condition mentioned with causal framing but not emphasized), 0.75 (condition presented as an important causal factor), 1.0 (condition presented as a primary or dominant cause). Models were instructed to output a JSON object with exactly eight keys matching the condition codes, each with a float value from the anchored scale. Temperature was set to 0.0 across all models to minimize stochastic variation within a model's responses. The ensemble design using four heterogeneous models serves a dual purpose: first, it distributes the extraction task across architectures with different training data and instruction-following tendencies, producing a more robust aggregate score; second, and critically, the variance across models on a given source-condition pair becomes a principled measure of epistemic uncertainty,

operationalized directly as the I component of the neutrosophic triplet. A model that reliably scores a condition at 0.8 across all four LLMs contributes low I; a condition where models range from 0.2 to 0.9 contributes high I regardless of the mean.

4.4 Structural baseline

Silencing analysis contrasts press T with a structural baseline derived from: (i) a validated perception survey (n=179 university students; conditions A, G, H via direct indicators); (ii) grey literature estimates (Crisis Group, OECO, Global Initiative; conditions B-F). Threshold for 'silenced': gap (press T - structural T) < -0.20.

The perception survey was administered to 179 university students in Guayaquil in 2024 as part of a broader assessment of community security perceptions (Batista Hernandez et al., 2023). Respondents rated the importance of gang fragmentation (condition A), institutional weakness (condition G), and economic pressure (condition H) as drivers of local violence on a five-point scale, which was converted to a [0,1] range using the same anchoring as the LLM scoring protocol. This provides a community-grounded baseline for those three conditions. For the remaining five conditions (B, C, D, E, F), structural T values were derived from grey literature estimates in criminological reports (Global Initiative, 2025; Crisis Group, 2025; OECO/PADF, 2024; UNODC, 2024). These estimates reflect the consensus of expert analytical organizations with field presence in Ecuador and were assigned T values corresponding to their qualitative characterization: conditions described as primary structural drivers received T = 0.75-0.80; those described as significant but secondary received T = 0.60-0.70. The mixed-baseline design introduces a heterogeneity that must be acknowledged when comparing silencing gaps across conditions (see Section 7).

4.5 N-fsQCA execution

We evaluated 41 configurations: 8 singletons, 28 pairs, 5 theoretically motivated triples. Each received a (T, I, F) triplet and an archetype classification.

The 41 configurations evaluated cover the complete set of singletons and all 28 pairwise intersections among the eight causal conditions. The five triples were selected on theoretical grounds: each represents a combination that criminological literature specifically identifies as potentially synergistic, such as the prison-drug route-gang fragmentation nexus that characterizes organized crime dynamics in Ecuador (Crisis Group, 2025; Atlantic Council, 2025). Evaluating the full 56-element space of triples was computationally feasible but would have introduced a multiple-comparisons problem without corresponding theoretical justification. For each configuration, the neutrosophic triplet (T, I, F) was computed using the formulas in Section 3.2, and the archetype was assigned according to the classification rules described there. Coverage — the proportion of outcome cases covered by the configuration — was computed $assum\big(min(X_T, Y_T)\big)/sum(Y_T)$, following Ragin (2008). All computations were performed in Python using the open-source pipeline available at github.com/mleyvaz/llm-nfsqca-narrative-violence.

5. Results

5.1 Corpus characterization

The corpus yields Y_violence_spike_T mean = 0.815 across LLMs, confirming universal topic focus. Local sources (n=16) are predominantly Ecuadorian media; international sources (n=15) span North American, European, and global analytical outlets.

The high mean outcome score (0.815) confirms that all sources are genuinely addressing violence escalation as their primary frame, validating the corpus selection criteria. The absence of sources with Y_T below 0.5 is a by-product of purposive corpus construction and is itself a limitation acknowledged in Section 7. Among local sources, Ecuadorian national dailies (El Universo, El Comercio) and digital investigative outlets (Primicias, EcuadorChequea) provide dense, event-specific coverage with strong causal framing, while regional outlets (Diario San Rafael, El Comun) contribute shorter descriptive accounts with weaker causal attribution. Among international sources, analytical organizations (Global Initiative, Crisis Group, InSight Crime, Atlantic Council) dominate in causal density, while wire service pieces (AP News, Al Jazeera) present briefer event-centered accounts with less structural framing. This heterogeneity in causal density within each group introduces variance that is partially captured by the I component of the N-fsQCA triplet.

5.2 N-fsQCA dominant configurations

Of 41 evaluated configurations, 40 achieve STRONG_SUFFICIENT (T >= 0.75, I <= 0.25). Table 1 shows the top 5.

Configuration	T	I	F	Coverage
D_territorial_war * F_prison_link	0.908	0.092	0.000	0.405
A_gang_fragmentation * D_territorial_war * F_prison_link	0.906	0.094	0.000	0.360
A_gang_fragmentation * F_prison_link	0.903	0.097	0.000	0.381
A_gang_fragmentation * D_territorial_war	0.898	0.102	0.000	0.585
B_drug_routes * D_territorial_war	0.883	0.117	0.000	0.499

Table 1. Top-5 N-fsQCA configurations. D=territorial war, F=prison linkages, A=gang fragmentation, B=drug routes.

The dominant configuration D*F (T=0.908, I=0.092) reflects criminological evidence that Ecuador's gangs are prison-based organizations directing territorial operations externally. F=0.000 across all configurations reflects corpus design (all articles address violence escalation); this limitation is addressed in §7.

The complementary configurations in Table 1 also merit interpretation. The three-way configuration A*D*F, combining gang fragmentation with territorial war and prison linkages (T = 0.906), achieves nearly identical consistency to the simpler D*F, which suggests that gang fragmentation is not an independent sufficient cause but rather a contextual elaboration of the dominant dyadic logic: prison-coordinated territorial warfare is the core mechanism, and gang fragmentation describes the organizational form through which that mechanism operates. The configuration A*D (T = 0.898, coverage = 0.585)

achieves the broadest coverage of any configuration in the table, meaning that it applies to the largest proportion of the outcome cases. This suggests that even without explicit mention of prison linkages, the combination of gang fragmentation and territorial war is the journalistic narrative structure most broadly shared across both local and international sources. The B*D configuration (T = 0.883, coverage = 0.499) signals the international-mediated framing: drug routes combined with territorial control is the explanatory dyad most characteristic of analytical organizations with a global trafficking perspective.

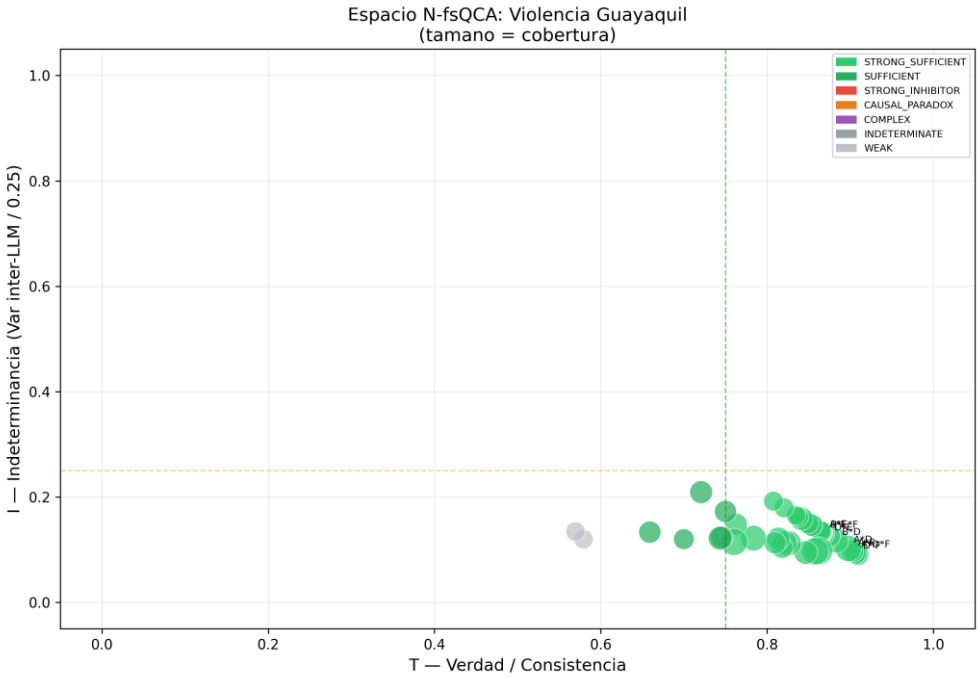

Figure 1. N-fsQCA space (T x I). Size = coverage. Green = STRONG_SUFFICIENT.

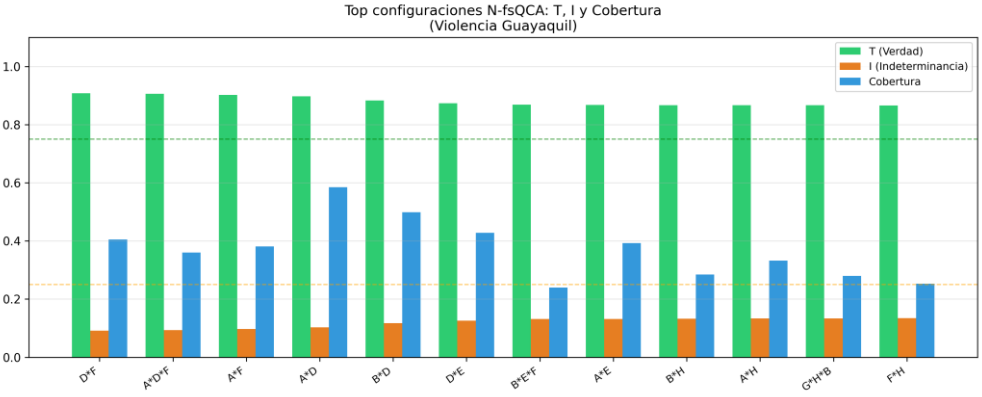

Figure 2. T, I, and coverage for top-12 configurations.

5.3 LLM model ablation

Table 2 reports per-model sufficiency consistency. All four models agree perfectly on weapons trafficking (1.000); economic pressure shows maximum disagreement (range 0.667-0.875), confirming that diffuse structural conditions are harder to extract from text than proximate/visible ones.

The pattern of cross-model agreement has an important methodological implication. Conditions that are lexically concrete and event-proximate in journalistic text, such as weapons trafficking and territorial war, tend to produce high inter-model agreement because the textual signals are direct: explicit references to firearms, zone control, or inter-gang combat leave little ambiguity for an LLM extraction system. Conditions that are structurally diffuse and expressed through indirect discourse, such as economic pressure and institutional weakness, require the model to infer a causal attribution from contextual framing rather than identify a discrete named entity or event type. The higher I values on these conditions therefore reflect a genuine property of the source texts, not a failure of the extraction system: when journalists gesture at structural causes without asserting them as primary drivers, the epistemic status of that causal attribution is legitimately uncertain. The I component is doing exactly the work it was designed to do.

Condition	Gemini-Flash	LLaMA-3.1	Phi-4	Qwen-3-8B
Weapons trafficking	1.000	1.000	1.000	1.000
Territorial war	1.000	1.000	0.958	0.960
Prison linkages	1.000	1.000	0.933	0.857
Extortion / vacunas	1.000	0.800	0.889	0.938
Gang fragmentation	0.947	0.938	1.000	0.947
Institutional weakness	0.900	0.846	0.893	0.857
Drug routes / ports	0.875	0.846	0.933	0.900
Economic pressure	0.667	0.818	0.875	0.833

Table 2. Per-LLM sufficiency consistency by condition.

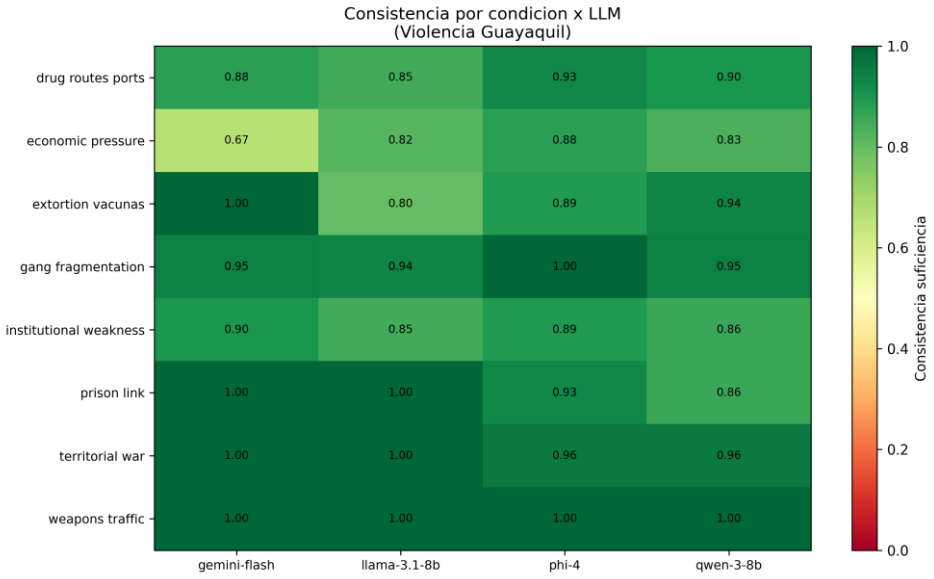

Figure 3. LLM consistency heatmap.

5.4 Local vs. international narrative framing

Table 3 compares mean T across media contexts. Two findings: (i) international media covers drug routes (gap=+0.330) and prison linkages (gap=+0.309) at roughly double the rate of local media; (ii) extortion/vacunas and weapons trafficking are covered equally by both contexts, indicating uncontested empirical recognition.

The drug routes and prison linkages asymmetry (gaps of +0.330 and +0.309 respectively) deserves particular attention. International analytical organizations invest substantially in mapping transnational trafficking networks and the role of prison gangs in coordinating them from within incarceration. Organizations such as InSight Crime, Global Initiative, and the Atlantic Council maintain field researchers and regional informant networks specifically oriented toward these structural dimensions. Local Ecuadorian media, operating under different commercial and security constraints, tend to report on immediate manifestations of violence: shootings, massacres, extortion demands, territorial confrontations. The structural architecture behind these events, which requires investigative resources and regional comparison, appears more prominently in international sources. This is not a failure of local journalism; it reflects a rational allocation of coverage resources given audience expectations and the proximity of threat. However, it means that readers of local media receive a systematically different causal theory of the crisis than readers of international analytical outlets, with direct implications for policy demands and electoral accountability.

Condition	Local T (ES)	Intl T (EN)	Gap (I-L)
Drug routes / ports	0.330	0.660	+0.330
Prison linkages	0.241	0.550	+0.309
Gang fragmentation	0.480	0.605	+0.125
Institutional weak.	0.519	0.617	+0.098
Economic pressure	0.322	0.377	+0.055
Territorial war	0.578	0.582	+0.004
Weapons trafficking	0.361	0.352	-0.009
Extortion / vacunas	0.387	0.340	-0.048

Table 3. Narrative framing by media context. Gap > +0.10: international substantially over-represents relative to local.

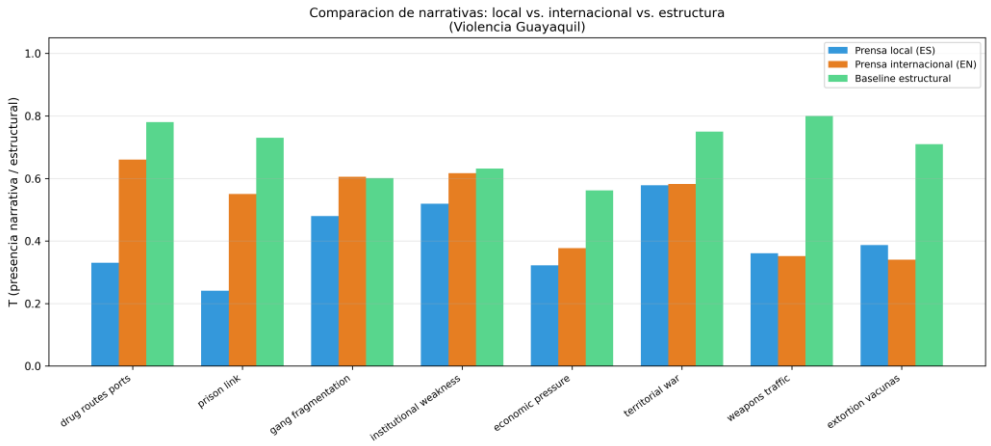

Figure 4. Causal framing: local (ES) vs. international (EN) vs. structural baseline.

5.5 Silencing analysis

Table 4 contrasts all-press T with the structural baseline. Five of eight drivers are systematically silenced (gap < -0.20). Weapons trafficking (-0.444) is most silenced despite near-perfect LLM agreement when mentioned, indicating editorial selection rather than extraction failure.

The weapons trafficking finding is particularly striking. Despite near-perfect cross-model agreement when the condition is mentioned ($T_L LM$ consensus = 1.000 in the ablation, indicating that all models recognize explicit weapons references identically), the overall press T for weapons trafficking is only 0.356. This dissociation means that the silence on weapons trafficking is not a measurement artifact: when journalists mention it, LLMs agree on its salience; but most sources simply do not mention it as a causal driver. The gap of -0.444 against the structural baseline ($T_s truct$ = 0.800) represents the largest silencing effect in the corpus. This is consistent with the political sensitivity of weapons trafficking in Ecuador: acknowledging the scale of firearms availability implicates international arms flows, domestic regulatory failures, and institutional complicity in ways that are commercially and legally risky for media organizations. Economic pressure shows a similar silencing dynamic (gap = -0.214), albeit against the lower survey-derived baseline rather than a higher grey literature estimate, suggesting that even the moderate causal attribution assigned to poverty and unemployment by community respondents exceeds what appears in either local or international press coverage.

Condition	Press T	Struct T	Gap	Silenced	Baseline
Weapons trafficking (E)	0.356	0.800	-0.444	YES	literature
Extortion / vacunas (C)	0.365	0.710	-0.345	YES	literature
Prison linkages (F)	0.390	0.730	-0.340	YES	literature
Drug routes / ports (B)	0.490	0.780	-0.290	YES	literature
Economic pressure (H)	0.348	0.562	-0.214	YES	survey n=179
Territorial war (D)	0.580	0.750	-0.170	no	literature
Institutional weakness (G)	0.566	0.632	-0.066	no	survey n=179
Gang fragmentation (A)	0.540	0.601	-0.061	no	survey n=179

Table 4. Silencing analysis. Threshold: gap < -0.20.

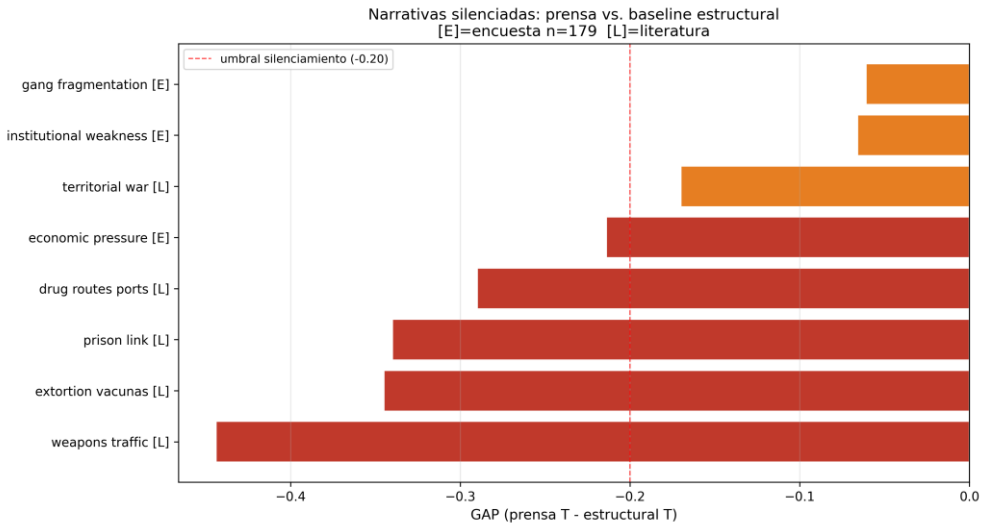

Figure 5. Silencing gaps: press T vs. structural baseline.

6. Discussion

6.1 Methodological implications

N-fsQCA extends fsQCA with two capabilities absent in the classical framework: (i) explicit quantification of epistemic indeterminacy via inter-annotator (inter-LLM) variance; (ii) detection of the paraconsistent regime (T + F > 1) in narrative corpora. The dominant configuration D*F achieves STRONG_SUFFICIENT status (T=0.908, I=0.092), indicating robust convergence that is not sensitive to the specific LLM used.

From a measurement theory perspective, the operationalization of I as inter-LLM variance is a principled rather than arbitrary choice. Inter-rater reliability is a standard concern in content analysis, typically handled by computing Cohen's kappa or Krippendorff's alpha and then discarding disagreement cases or averaging over raters. The N-fsQCA approach instead preserves disagreement as structural information. A condition with low inter-LLM agreement is not an extraction failure to be corrected; it is a signal that the causal attribution in the source text is genuinely ambiguous, and that ambiguity has substantive meaning. This reframing transforms what is typically treated as measurement noise into a first-class theoretical variable. The approach generalizes beyond LLM scoring: any setting where multiple independent coders assign condition memberships can operationalize I as coder variance, extending N-fsQCA to human annotation studies with full backward compatibility to classical fsQCA when I is universally low.

6.2 Geopolitics of causal framing

The local/international framing divergence (Section 5.4) operationalizes a theoretical claim from decolonial media studies: international outlets frame Guayaquil's crisis through global supply-chain and prison-governance lenses, while local media foregrounds lived community experience (extortion, gang fragmentation). Neither frame is incorrect; they imply different policy responses.

The policy implications are direct. If security policy in Ecuador is primarily informed by local media, the policy agenda will be weighted toward reactive and kinetic responses: gang suppression, territorial containment, visible policing. If it is informed by international analytical coverage, the emphasis will shift toward port interdiction, prison reform, and transnational law enforcement cooperation. If it is informed by community perceptions as captured in the survey baseline, the priority will shift again toward economic pressure, weapons control, and institutional accountability. These are not three descriptions of the same policy; they are three genuinely different theories of change with different institutional constituencies, different time horizons, and different distributional consequences for Guayaquil residents. N-fsQCA provides a framework for making these divergent theories of change explicit, comparable, and subject to empirical scrutiny rather than implicit and resistant to challenge.

6.3 Decolonial critique operationalized

The silencing analysis (Section 5.5) provides quantitative evidence for the decolonial critique: structural drivers recognized in community perceptions (survey, n=179) are systematically under-represented in both local and international press. The finding that economic pressure (gap=-0.214) is silenced only when the survey-derived baseline is used — rather than higher literature estimates — illustrates the methodological importance of empirical baselines over authorial judgment.

The methodological innovation in the silencing analysis is the use of a community-grounded empirical baseline rather than a normative claim about what should be covered. By anchoring three conditions in a validated perception survey of n = 179 Guayaquil residents, the analysis avoids the circularity of measuring silencing against a researcher's prior about what matters. The finding that economic pressure is silenced only when measured against the survey baseline, rather than against higher grey literature estimates, is itself informative: it suggests that the structural literature overestimates the salience of economic pressure as a proximate causal driver relative to community experience, or conversely that community respondents are more focused on lived economic conditions than structural analysts who emphasize political economy. Resolving this tension requires richer qualitative work with affected communities, which N-fsQCA can motivate and orient but not replace.

7. Limitations

Outcome selection bias.

All 31 sources address violence escalation (Y_mean = 0.815), producing F = 0 universally. This design choice was deliberate: the research question concerns how sources explain a specific shared event, not whether different source types differ in their violence recognition. However, it means the corpus cannot detect the paraconsistent regime (T + F > 1) that is theoretically possible in N-fsQCA. A study comparing coverage of a high-violence period against a low-violence counterfactual would allow F > 0 to emerge for conditions whose causal role is contested across time periods, enabling a richer test of the paraconsistent detection capability of the framework. Future work should construct such temporally stratified corpora.

All 31 sources address violence escalation (Y_mean=0.815), producing F=0 universally. Future work must include low-violence comparison periods.

Survey baseline scope.

The perception survey (n = 179 university students) covers only three of eight structural conditions (A, G, H) and uses a convenience sample of tertiary-educated respondents in a specific urban context. University students are not representative of the communities most affected by violence in Guayaquil, which tend to be peri-urban settlements with lower educational attainment and higher direct exposure to gang activity. The survey-derived baselines for gang fragmentation, institutional weakness, and economic pressure should therefore be interpreted as the perceptions of a relatively privileged urban subsample rather than a community-wide baseline. Future studies should deploy stratified surveys across affected neighborhoods, ideally using validated instruments from the criminological literature on community violence perception.

The perception survey (n=179 university students) covers only 3 of 8 structural conditions (A, G, H) and uses a convenience sample. Community-stratified surveys are needed.

Literature baselines.

Five structural T values (B through F) derive from grey literature estimates in criminological reports rather than quantitative data. These estimates introduce author judgment at the baseline-construction stage, creating the risk that the silencing analysis reflects researcher priors rather than empirical structural conditions. Official crime statistics from UNODC (2024) and Ecuador's Ministry of Government provide some grounding, but their granularity at the condition level is limited: homicide counts and trafficking seizures do not directly translate into causal salience scores on a [0,1] scale. Replacing grey literature baselines with survey-derived measures for all eight conditions, using community-stratified samples, would substantially strengthen the silencing analysis and reduce the interpretive ambiguity introduced by heterogeneous baseline sources.

Five structural T values (B-F) derive from grey literature, introducing author judgment. Official crime statistics (UNODC, Ministry of Government Ecuador) should replace these.

LLM reliability.

Mistral-7B v0.1 was excluded from the final ensemble after producing malformed JSON output in 7 of 31 scoring attempts. This exclusion is itself a finding: model reliability on structured extraction tasks is not uniform across the open-source ecosystem, and the decision to exclude or include a model affects both the ensemble mean T and the variance-based I. The four retained models were selected partly on the basis of their JSON compliance rates in preliminary testing, which introduces a selection bias toward models that are better at instruction-following rather than necessarily better at causal understanding. Robustness checks with larger and more capable models, including proprietary systems such as GPT-4o and Claude 3.7, would clarify whether the patterns observed here generalize across the capability spectrum or are specific to the open-source models used.

Mistral-7B was excluded after JSON failures in 7/31 cases. Model reliability is itself a methodological variable requiring explicit management.

Causal status.

The pipeline analyzes narrated causality rather than actual causal mechanisms. A high T value on a configuration means that journalistic sources consistently present that configuration as causally relevant, not that the configuration actually drives violence outcomes in the real world. The distinction is methodologically important: media causal theories can be systematically wrong, omitting structural causes (as the silencing analysis demonstrates), or systematically biased toward visible and proximate causes at the expense of diffuse structural ones. N-fsQCA provides a rigorous description of journalistic causal epistemology, not a causal model of violence. Connecting narrated causality to empirical causal mechanisms requires additional causal inference work, such as difference-in-differences designs exploiting natural experiments in coverage variation, which is beyond the scope of the present paper.

The pipeline analyzes narrated causality, not actual causal mechanisms. High T reflects journalistic co-narration, not empirical causation.

8. Conclusion

This paper introduces N-fsQCA, a neutrosophic extension of fsQCA, and applies it to a bilingual corpus of 31 sources on Guayaquil's violence crisis. Three findings emerge: (i) territorial war combined with prison linkages is the most consistent causal configuration (T=0.908); (ii) international media foregrounds drug routes and prison links at double the rate of local media, reflecting a geopolitical framing asymmetry; (iii) five structural drivers are systematically silenced in the press. The pipeline is open-source and replicable at < USD 5.

Future work should extend N-fsQCA in three directions. First, temporal extension: applying the pipeline to a longitudinal corpus spanning the full 2020-2026 period would allow tracking of how journalistic causal theories evolve in response to policy interventions, criminal organization restructuring, and electoral cycles. Second, outcome variation: incorporating sources from low-violence periods would activate the F component and enable detection of the paraconsistent regime. Third, community integration: pairing the media corpus analysis with deeper community surveys and participatory research would ground the silencing analysis in lived experience rather than expert estimation, strengthening the decolonial critique and providing actionable guidance for both journalists and policymakers working on urban security in Latin America.

Beyond the empirical findings, the paper makes a methodological argument: that indeterminacy is not noise to be minimized but signal to be preserved. The N-fsQCA framework transforms the standard problem of inter-annotator disagreement from a quality-control challenge into a substantive variable. When four heterogeneous language models disagree on the causal salience of economic pressure in a news article, that disagreement reflects genuine textual ambiguity, which in turn reflects the way journalists hedge, qualify, and contextualize structural causes in ways that resist clean binary attribution. Capturing this hedging as I rather than discarding it as noise is epistemically more honest and analytically more informative.

Reproducibility Statement

Code, data, and figures: github.com/mleyvaz/llm-nfsqca-narrative-violence (MIT License). Requires OpenRouter API key (< USD 5 for full corpus).

CRediT Author Statement

Maikel Leyva-Vazquez: Conceptualization, Methodology, Software, Formal analysis, Writing (original draft), Visualization, Project administration. Noel Batista Hernandez: Conceptualization (survey instrument), Data curation (survey n=179), Writing (review & editing), Validation. Florentin Smarandache: Conceptualization (Neutrosophic framework), Writing (review & editing), Supervision.

Conflict of Interest

M. Leyva-Vazquez is Editor-in-Chief of Neutrosophic Sets and Systems. F. Smarandache is founder of neutrosophic logic and associated with that journal. This work was not submitted to that journal. No other conflicts declared.

References

ACLED (2025). Ecuador's Noboa declared war on 22 gangs. ACLED Report. https://acleddata.com

Atlantic Council (2025). Cradle of violence: How criminal networks are winning Ecuador's youth. https://www.atlanticcouncil.org

Batista Hernandez, N., Leyva-Vazquez, M., & Smarandache, F. (2023). Neutrosophic methods in educational evaluation. Neutrosophic Sets and Systems, 58, 45-62.

Birhane, A. (2021). Algorithmic injustice: a relational ethics approach. Patterns, 2(2), 100205.

Crisis Group (2025). Paradise Lost? Ecuador's Battle with Organised Crime. Report No. 109. https://www.crisisgroup.org

Da Costa, N. C. A. (1974). On the theory of inconsistent formal systems. Notre Dame Journal of Formal Logic, 15(4), 497-510.

Dusa, A. (2019). QCA with R. Springer.

Farquhar, S., Kossen, J., Kuhn, L., & Gal, Y. (2024). Detecting hallucinations in large language models using semantic entropy. Nature, 630(8017), 625-630. https://doi.org/10.1038/s41586-024-07421-0

Global Initiative (2025). Ecuador's criminal crisis. https://globalinitiative.net

HRW (2026). World Report 2026: Ecuador. Human Rights Watch.

InSight Crime (2025a). Massacres, Alliances, and Betrayals: Inside the Gang War That Tore Ecuador's Tiguerones Apart.

InSight Crime (2025b). GameChangers 2024: Ecuador Finds Victory Elusive in War on Gangs.

InSight Crime (2025c). Why Have Ecuador's Drug Trafficking Gangs Turned to Extortion?

Leyva-Vazquez, M. Y., & Smarandache, F. (2026, in prep.). The Third Answer. NSIA Publishing.

Manakul, P., Liusie, A., & Gales, M. J. F. (2023). SelfCheckGPT. EMNLP.

Mhlambi, S. (2020). From rationality to relationality. Carr Center Discussion Paper 2020-009.

Misangyi, V. F. et al. (2017). Embracing causal complexity. Journal of Management, 43(1), 255-282.

OECO / PADF (2024). Radiografia de la extorsion en Ecuador.

Ragin, C. C. (1987). The comparative method. University of California Press.

Ragin, C. C. (2008). Redesigning social inquiry. University of Chicago Press.

Smarandache, F. (1998). Neutrosophy. American Research Press.

Smarandache, F. (2003). Definition of neutrosophic logic. Proceedings of the First International Conference on Neutrosophy.

Tonmoy, S. M. T. I., Zaman, S., Jain, V., Rani, A., Rawte, V., Chadha, A., & Das, A. (2024). A comprehensive survey of hallucination mitigation techniques in large language models. arXiv preprint arXiv:2401.01313. https://arxiv.org/abs/2401.01313

UNODC (2024). Global study on homicide. United Nations.

Wang, H., Smarandache, F., Zhang, Y., & Sunderraman, R. (2010). Single valued neutrosophic sets. Multispace and Multistructure, 4, 410-413.

Fairclough, N. (1995). Media discourse. Edward Arnold.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), Syntax and Semantics, Vol. 3: Speech Acts (pp. 41-58). Academic Press.

Hobbs, J. R. (1990). Literature and cognition. CSLI Lecture Notes, Vol. 21. Stanford: Center for the Study of Language and Information.

Levinson, S. C. (1983). Pragmatics. Cambridge University Press.

Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. Text, 8(3), 243-281.

Sperber, D., & Wilson, D. (1995). Relevance: Communication and cognition (2nd ed.). Blackwell.

Van Dijk, T. A. (2006). Discourse and manipulation. Discourse & Society, 17(3), 359-383.